

07-07-00

A

07/06/00
JC845 U.S. PTO

JC857 U.S. PTO
09/610278
07/06/00

Attorney Docket No. IBN.5203

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

BOX PATENT APPLICATION
Assistant Commissioner for Patents
Washington, D.C. 20231

NEW APPLICATION TRANSMITTAL

Transmitted herewith for filing is the patent application of

Inventor(s): Ross W. Callon, *et al.*

For (title): APPARATUS AND METHOD FOR
EFFICIENT HASHING IN
NETWORKS

1. **Type of Application** This new application is for
a(n)

- ☒ Original (nonprovisional)
☐ Design
 ☐ Plant
☐ Divisional.
☐ Continuation.
☐ Continuation-in-part (C-I-P).

2. **Papers Enclosed**

11 Pages of specification
5 Pages of claims
1 Page of Abstract
2 Sheets of drawings
 ☐ formal
 ☒ informal
1 Page of Cover Sheet

☐ The enclosed drawing(s) are photograph(s), and there is also attached a "PETITION TO ACCEPT PHOTOGRAPH(S) AS DRAWING(S)." 37 C.F.R. 1.84(b).

CERTIFICATE OF MAILING

37 C.F.R. § 1.10

"Express Mail" Mailing Label Number EL584588998US
I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to BOX PATENT APPLICATION, Assistant Commissioner for Patents, Washington, DC 20231.

Date: 7/6/00

Vanessa Marakas
Vanessa Marakas

Vanessa Marakas
Print Name

3. Additional papers enclosed

- ☐ Preliminary Amendment
- ☐ Information Disclosure Statement (37 C.F.R. 1.98)
- ☐ Form PTO-1449 (PTO/SB/08A and 08B)
- ☐ Copies of cited references
- ☐ Declaration of Biological Deposit
- ☐ Submission of "Sequence Listing," computer readable copy and/or amendment pertaining thereto for biotechnology invention containing nucleotide and/or amino acid sequence.
- ☐ Authorization of Attorney(s) to Accept and Follow Instructions from Representative
- ☐ Special Comments
- ☒ Other: Return Postcard.

4. Declaration or oath

- ☒ Enclosed
 - ☒ Unexecuted
 - ☐ Executed by
 - ☐ inventors
 - ☐ legal representative of inventor(s).
37 CFR 1.42 or 1.43.
 - ☐ joint inventor or person showing a proprietary interest on behalf of inventor who refused to sign or cannot be reached.
 - ☐ This is the petition required by 37 CFR 1.47 and the statement required by 37 CFR 1.47 is also attached. See item 12 below for fee.
- ☐ Not Enclosed
- ☐ Application is made by a person authorized under 37 C.F.R. 1.41 (c) on behalf of all the above named inventor(s).
 - ☐ Showing that the filing is authorized.

5. Assignment

- ☒ An assignment of the invention to IronBridge Networks, Inc.
 - ☐ is attached. A separate ☐ "COVER SHEET FOR ASSIGNMENT (DOCUMENT) ACCOMPANYING NEW PATENT APPLICATION" or ☒ FORM PTO 1595 is also attached.
 - ☒ will follow.

6. Certified Copy

Certified copy(ies) of application(s)

Country	Appln. no.	Filed
---------	------------	-------

Country	Appln. no.	Filed
---------	------------	-------

Country	Appln. no.	Filed
---------	------------	-------

from which priority is claimed

- ☐ is (are) attached.
☐ will follow.

7. Fee Calculation (37 C.F.R. 1.16)

CLAIMS AS FILED

Number filed	Number Extra	Rate	Basic Fee 37 C.F.R. 1.16(a) <u>\$690.00</u>
Total			
Claims (37 CFR 1.16(c)) 28	- 20 = 8	\$ 18.00	\$144.00
Independent			
Claims (37 CFR 1.16(b)) 2	- 3 = 0	\$ 78.00	
Multiple dependent claim(s), if any (37 CFR 1.16(d))			

- ☐ Amendment cancelling extra claims is enclosed.
☐ Amendment deleting multiple-dependencies is enclosed.
☐ Fee for extra claims is not being paid at this time.

Filing Fee Calculation

\$834.00

8. Small Entity Statement(s)

☒ Verified Statement(s) that this is a filing by a small entity under 37 CFR 1.9 and 1.27 is (are) attached.

☐ Status as a small entity was claimed in prior application _____, filed on _____ from which benefit is being claimed for this application under:

35 U.S.C. ☐ 119(e),

☐ 120,

☐ 121,

☐ 365(c),

and which status as a small entity is still proper and desired.

☐ A copy of the verified statement in the prior application is included.

Filing Fee Calculation (50% of A, B or C above)

\$ 417.00

9. Fee Payment Being Made at This Time

☐ Not Enclosed

☐ No filing fee is to be paid at this time.

(This and the surcharge required by 37 C.F.R. 1.16(e) can be paid subsequently.)

☒ Enclosed

☒ Basic filing fee

\$ 417.00

☐ Recording assignment

(\$40.00; 37 C.F.R. 1.21(h))

(See attached "COVER SHEET FOR ASSIGNMENT
ACCOMPANYING NEW APPLICATION".)

\$ _____

Total fees enclosed

\$ 417.00

10. **Method of Payment of Fees**

☒ Check in the amount of \$ 417.00

☐ Charge Account No. 19-0079 in the amount of \$ _____
A duplicate of this transmittal is attached.

11. **Authorization to Charge Additional Fees**

☒ The Commissioner is hereby authorized to charge the following additional fees during the entire pendency of this application to Account No. 19-0079.

☒ 37 C.F.R. 1.16(a), (f) or (g) (filing fees)

☒ 37 C.F.R. 1.16(b), (c) and (d) (presentation of extra claims)

☐ 37 C.F.R. 1.16(e) (surcharge for filing the basic filing fee and/or declaration on a date later than the filing date of the application)

☐ 37 C.F.R. 1.17 (application processing fees)

☐ 37 C.F.R. 1.18 (issue fee at or before mailing of Notice of Allowance, pursuant to 37 C.F.R. 1.311(b))


12. **Instructions as to Overpayment**

☒ Credit Account No. 19-0079

☐ Refund

Date: July 6, 2010
Samuels, Gauthier & Stevens, LLP
225 Franklin Street, Suite 3300
Boston, MA 02110
Telephone: (617) 426-9180, Ext. 149
Facsimile: (617) 426-2275

Respectfully submitted,


Steven M. Mills
Registration Number 36,610
Attorney for Applicants

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant(s): Ross W. Callon, *et al.*
Filing Date: July 5, 2000
Title: APPARATUS AND METHOD FOR EFFICIENT HASHING IN NETWORKS

CERTIFICATE OF MAILING UNDER 37 C.F.R. § 1.10

"Express Mail" Mailing Label Number EL584588998US I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated below and is addressed to BOX PATENT APPLICATION, Assistant Commissioner for Patents, Washington, DC 20231.

July 6, 2000
Date

Vanessa Marakas
Vanessa Marakas

BOX PATENT APPLICATION

Assistant Commissioner for Patents
Washington, DC 20231

TRANSMITTAL LETTER

Sir:

Enclosed herewith for filing in the above-identified patent application please find the following listed items:

1. New Application Transmittal;
2. New Patent Application;
3. Unexecuted Declaration, Petition and Power of Attorney;
4. Two (2) Sheets of Informal Drawings;
5. Check in the amount of \$417 to cover filing fee;
6. Unexecuted Verified Statement Claiming Small Entity Status; and
7. Return Postcard.

In connection with the foregoing matter, please charge any additional fees which may be due, or credit any overpayment, to Deposit Account Number 19-0079. A duplicate copy of this letter is provided for this purpose.

Respectfully submitted,

Date: July 6, 2000
Samuels, Gauthier & Stevens, LLP
225 Franklin Street, Suite 3300
Boston, MA 02110
Telephone: (617) 426-9180, Ext. 149
Facsimile: (617) 426-2275
K:\IronBridge\5203\transapp.wpd

Steven M. Mills
Steven M. Mills
Registration Number 36,610
Attorney for Applicants

09640278-070600

Applicant or Patentee: Ross W. Callon et al. Attorney's
Serial or Patent Number: N/A Docket No: IBN.5203
Filed or Issued: Filed herewith
For: Apparatus and Method for Efficient Hashing in Networks

**VERIFIED STATEMENT (DECLARATION) CLAIMING SMALL ENTITY
STATUS (37 CFR 1.9 (f) and 1.27 (b)) -- SMALL BUSINESS CONCERN**

I hereby declare that I am

- ☐ the owner of the small business concern identified below:
☒ an official of the small business concern empowered to act on behalf of the concern identified below:

NAME OF CONCERN IronBridge Networks, Inc.
ADDRESS OF CONCERN 55 Hayden Avenue, Lexington, Massachusetts 02173

I hereby declare that the above identified small business concern qualifies as a small business concern as defined in 13 CFR 121.3-18, and reproduced in 37 CFR 1.9(d), for purposes of paying reduced fees under section 41(a) and (b) of Title 35, United States Code, in that the number of employees of the concern, including those of its affiliates, does not exceed 500 persons. For purposes of this statement, (1) the number of employees of the business concern is the average over the previous fiscal year of the concern of the persons employed on a full-time, part-time or temporary basis during each of the pay periods of the fiscal year, and (2) concerns are affiliates of each other when either, directly or indirectly, one concern controls or has the power to control the other, or a third party or parties controls or has the power to control both.

I hereby declare that rights under contract or law have been conveyed to and remain with the small business concern identified above with regard to the invention, entitled:

Apparatus and Method for Efficient Hashing in Networks

Ross W. Callon et al. by inventor(s)
described in

- ☒ the specification filed herewith
☐ application serial no. _____, filed _____
☐ patent no. _____, issued _____

If the rights held by the above identified small business concern are not exclusive, each individual, concern or organization having rights to the invention is listed below* and no rights to the invention are held by any person, other than the inventor, who could not qualify as a small business concern under 37 CFR 1.9 (d) or by any concern which would not qualify as a small business concern under 37 CFR 1.9 (d) or a nonprofit organization under 37 CFR 1.9 (e).

*NOTE: Separate verified statements are required from each named person, concern or organization having rights to the invention averring to their status as small entities. (37 CFR 1.27)

FULL NAME NONE
ADDRESS _____

☐ INDIVIDUAL ☐ SMALL BUSINESS CONCERN ☐ NONPROFIT ORGANIZATION

FULL NAME _____
ADDRESS _____

☐ INDIVIDUAL ☐ SMALL BUSINESS CONCERN ☐ NONPROFIT ORGANIZATION

I acknowledge the duty to file, in this application or patent, notification of any change in status resulting in loss of entitlement to small entity status prior to paying, or at the time paying, the earliest of the issue fee or any maintenance fee due after the date on which status as a small entity is no longer appropriate. (37 CFR 1.28 (b))

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment or both, under section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application, any patent issuing thereon, or any patent to which this verified statement is directed.

NAME OF PERSON SIGNING Kathleen Huber
TITLE OF PERSON OTHER THAN THE OWNER Vice President
ADDRESS OF PERSON SIGNING 55 Hayden Avenue, Lexington, Massachusetts, 02173

SIGNATURE _____

DATE _____

Attorney Docket Number IBN.5203

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

PATENT APPLICATION

of

Ross W. Callon

James D. Carlson

Douglas C. Ehlert

and

Yuval Peduel

for

APPARATUS AND METHOD FOR EFFICIENT HASHING IN NETWORKS

09640278-070600

APPARATUS AND METHOD FOR EFFICIENT HASHING IN NETWORKS

Field of the Invention

The present invention is related to techniques for forwarding data packets on networks and is more particularly related to a technique for forwarding encapsulated data packets on a network in which network nodes are in general connected by multiple paths.

Background of the Invention

In digital computer networks such as the Internet, collections of data, referred to as "datagrams", are typically transferred from node to node over the network in packets. Each packet of data typically includes a header portion and a data portion. In accordance with the common Internet protocol (IP), the header portion typically includes a 32-bit source identifying portion which identifies the source node that originated the packet and a 32-bit destination identifying portion which identifies the destination node to which the packet is ultimately to be transferred. A packet header can also include other predefined fields such as a protocol field.

At each node, a router is used to forward the packet to the next node in the path toward the destination node. When a router receives a packet, it examines the destination address in the packet header. It then searches its locally stored routing table to determine the next node to which the packet should be transferred in order to ensure that it will reach its destination, typically along the shortest possible path. The router then forwards the packet to the next node identified in the routing table. This process continues at each successive node until the destination node is reached.

In many cases, in such a datagram IP network, when forwarding an IP packet, there are situations in which there are two or more choices for the next step or "hop" that the packet can take, i.e., the next node to which the packet is to be forwarded. FIG. 1

contains a schematic block diagram of a conventional IP packet forwarding network 10. The network 10 includes multiple nodes 12 connected by links 13. Referring to FIG. 1, the case in which IP packets are forwarded from node A to node F, for example, is considered. In this situation, node A will forward the packet to node B. Node B can forward the packet toward node E along a first path via node C or along a second path via node D.

In general, multiple hosts 14 are coupled to each node A and F. A host 14 coupled to node A may have a sequence of multiple IP packets destined for another host 14 attached to node F. It is desirable to keep the packets associated with any one host-to-host flow in order. This is important to improve the efficiency of communication. For example, in many cases, the hosts 14 may be running Internet applications over the Transmission Control Protocol (TCP), and TCP may make use of "slow start." When applications are making use of TCP slow start, if packets are delivered out of order, the TCP implementation assumes that the misordering of packets is caused by congestion in the network. In response, the rate of traffic transmitted may be reduced. If in fact there is no congestion in the network, then this will result in less efficient use of the network. In any case, it is often advisable if not critical that packets associated with the same source/destination node pair be transferred on the same series of links.

Typically, IP routers solve this problem by choosing between multiple equal-cost choices for the next hop for a particular packet. The router typically performs an analysis of the packet header contents to assign each packet to a link. Usually, this involves a hash function of the five-tuple of fields in the IP header, i.e., source IP address, destination IP address, protocol, source port number, destination port number, or a subset of these fields, such as source IP address and destination IP address. A hash function is designed to perform a computation on one or more data words and return a unique data word of shorter length. For example, a hash function performed on two 32-bit IP addresses may divide the combined 64-bit word by a constant and return as a

result the value of the reminder in fewer bits, e.g., five. Other hash procedures include the use of a cyclic redundancy check (CRC) and/or the use of a checksum.

Each time a hash procedure is performed on the same initial values, the same result is obtained. This ensures that packets associated with any one source/destination pair always take the same path, while simultaneously allowing different packets to take different paths. As noted above, sometimes additional fields may be used for the hash. It is noted that any packets belonging to the same flow of packets, i.e., packets which should be kept in order, will also contain the same protocol field in the IP header. Packets which contain a different value in the protocol field may therefore be safely transmitted on a different path. Similarly, if the protocol field indicates that the next higher level protocol is TCP, then packets which contain different TCP port numbers can be routed on different paths. For these reasons, it is common for the hash to also take account of the protocol and port fields.

Thus, referring to FIG. 1, in general, multiple hosts 14 attached to node A send IP packets to multiple hosts 14 attached to router F. Under the technique described above, packets from any one source/destination pair will always be transmitted over the same path. Node B performs the hash function on each incoming packet. The result of the hash function is used to select either node C or node D as the next node, i.e., the node to which the packet is to be forwarded. Because the hash operation is performed on information in the packet that uniquely identifies its associated source/destination pair, all packets associated with that pair will be forwarded by node B to the same node, i.e., along the same path. The packets averaged over all of the source/destination pairs will in general be split, with some being sent along the path that includes node C and some being sent along the path that includes node D. This allows more efficient loading of the network 10 by splitting traffic among multiple available paths.

As the demand for data network services increases, it is becoming increasingly common for the interconnection between any two nodes to include multiple parallel

links. Using multiple links increases the total bandwidth available for data transmission. Also, using multiple links allows for the possibility that if one link fails, there will still be a path through the network between any two nodes.

FIG. 2 is a schematic block diagram of a network 100 which includes multiple links 113 between nodes 112. Specifically, the nodes B, C, D and E in the core of the network 100 are shown interconnected using two parallel links 113 rather than a single link.

In this case, the same technique as described above for forwarding packets can be used. In particular, node B can perform a hash on the IP source and destination addresses. In this case, node B has four choices for possible links to use in forwarding a packet toward node F. Node B can therefore use a hash function with four possible output values. Each of the four links is considered a possible choice for forwarding of the data packet. In this case, as in the previous case, packets for any single source/destination pair will always go via the same link, i.e., via either one of the two links to C or one of the two links to D. However, the packets averaged over all of the source/destination pairs will generally be split, with some being sent via each of the four links. This allows for more efficient loading of the network by splitting traffic among multiple available paths as well as multiple available links within a path.

A problem occurs in conventional hashing which substantially reduces the effectiveness of the function. In a typical network, the same hashing function is used at every node in the network. Consequently, for a particular packet being forwarded along a path, the same hash function result is obtained at every node. Also, the packets that actually arrive at a particular node in the network are a function of the hash values that have been computed at other upstream nodes. This implies that the hash function, as implemented at any particular node, will not obtain all possible output values with the same probability. As a result of these factors, the conventional hashing function does not provide for utilization of the network resources in the most efficient fashion.

Summary of the Invention

5 The present invention provides an approach to hashing which effectively performs a different hashing operation at each node along a path, such that the above drawbacks of prior approaches are avoided. The invention is directed to an apparatus and method for transferring data on a network which includes a plurality of nodes connected by a plurality of links. At least one of the nodes is connected to multiple links over which data can be forwarded from the node onto the network toward another node. The data includes an identifying portion which identifies a source node which originated the data and a destination node to which the data is to be transferred. The at least one node is associated with a node identifying value which serves to identify the node and distinguish the node from other nodes on the network. A link selection value, which is used to identify the one of the links out of the node on which the data is to be transferred, is generated using the identifying portion of the data and the node identifying value. Because the selection value is generated using the node identifying value, which is unique to the particular node, the selection value will depend on the particular node generating the value and will be different at each node, even for the same data having the same data identifying portion.

10 15 The data can be a packet of data and, in particular, can be an Internet Protocol (IP) data packet. The identifying portion of the packet can be at least a portion of the IP packet header. For example, the identifying portion can include the source node ID, the destination node ID and/or the protocol portion of the IP packet header. The node identifying value can be a single constant value which is configured in each router or can be a portion, e.g., the low-order bits, of the IP address of the node router.

20 25 In one embodiment, in generating the link selection value, a logical operation, for example, a hash operation, is performed on the identifying portion of the data and the node identifying value. Again, the identifying portion of the data can include the source ID, the destination ID and/or the protocol fields of the header portion of the data packet.

The hash operation can include a cyclic redundancy check (CRC) performed on the identifying portion of the data and the node identifying value.

Brief Description of the Drawings

5 The foregoing and other objects, features, and advantages of the invention will be apparent from the following more particular description of preferred embodiments of the invention, as illustrated in the accompanying drawings in which like reference characters refer to the same parts throughout the different views. The drawings are not necessarily to scale, emphasis instead being placed upon illustrating the principles of the invention.

10 FIG. 1 is a schematic block diagram of a network having multiple paths between nodes.

FIG. 2 is a schematic block diagram of a network having plural links between nodes.

15 FIG. 3 is a schematic block diagram of a network having multiple paths between nodes and one path with plural links between nodes.

FIG. 4 is a schematic block diagram of a network having multiple paths between nodes.

Detailed Description of Preferred Embodiments

20 FIG. 3 is a schematic block diagram of a configuration of a network 200 which is subject to one of the drawbacks of conventional hash functions. The network 200 is similar to the network 10 of FIG. 1, except that the network 200 uses a pair of parallel links 213a and 213b between nodes C and E. Here, in general, there are multiple packets being transmitted from hosts 14 which are attached directly or indirectly to router A toward hosts 14 which are attached directly or indirectly to router F.

25 In this case, router B may use multi-path, sending some packets via router C and some via router D. In a conventional hashing configuration, all of the routers A-F are

implemented with the same hashing function. Given that the packets have various source addresses and/or various destination addresses, it is possible for router B to perform a hash function on the source and destination addresses in order to send some packets via node C and some via node D while sending all packets corresponding to a particular source/destination pair via the same path. This ensures that a good split of traffic is achieved on the B-C and B-D links, while preserving packet order within each individual host-to-host flow.

In the network 200 of FIG. 3, router C will forward all packets to node F via node E. However, it is noted that node C has two links to node E. Therefore, it is desirable to split the traffic equally across the two links. However, in most situations, router C will be an identical implementation to router B, or at least will implement the same hashing function to select paths for forwarding packets. As a result, router C will have the same output of the hashing function.

To illustrate the problem, it is assumed that the hashing function in the routers in the network 200 consists of a cyclic redundancy check (CRC) function on the IP source and destination addresses of packets. It is also assumed that all packets with an even result are sent on one link, and all packets with an odd result are sent on the other link. For example, at node B, it is assumed that all packets with an even result are sent to node C, and all packets with an odd result are sent to node D. Also, at node C, all packets with an even result of the CRC are sent on the upper link 213a to E, and all packets with an odd result are sent on the lower link 213b to node E.

The problem here is that all of the packets that arrive at node C (from A to F) will be packets for which the CRC hash function in this example returns an even result. Given that node C will perform the exact same hash function on the exact same fields in the IP packet, the same result will occur, meaning that all packets will result in an even hash value when the function is evaluated at node C. As a result, all packets will take

the upper link 213a to node E, and the lower link 213b will not be used at all. This is a very unequal and inefficient utilization of the links from node C to E.

FIG. 4 is a schematic block diagram of another configuration of a network 300 which experiences a drawback in connection with a conventional hashing approach. In the network 300 illustrated in FIG. 4, data traffic is forwarded from hosts 14 attached to node A to hosts 14 attached to node G. Node C uses a hash function to select one of the two links 313c and 313f used to forward data toward node G via node F or node E, respectively. For illustration, it is again assumed that node B sends packets with an even hash value to node C and packets with an odd hash value to node D. It is also assumed that node C sends packets with an even hash value to node E and packets with an odd hash value to node F.

In this case, because of the hash operation performed at node B, node C receives only packets with even hash values. When node C performs its hash operation on the packets, only even results will be obtained, and all packets will be forwarded on link 313f to node E and on link 313g to node G. The links 313c to node F and 313h to node G would not be used at all, again resulting in very inefficient use of network resources.

In accordance with the invention, a different hash operation is performed at each node along a path such that these problems in conventional hashing approaches are eliminated. This approach ensures that a relatively equal splitting of traffic can be achieved at each node such that network resources are more efficiently utilized.

For general applications to networks, it would be difficult to implement a different hashing procedure at each node on a network, given the large number of nodes. Instead, the hashing operation of the invention operates on both specific fields in the data packet and an additional input value that is unique to each router. For example, the additional input value can be a single constant value configured in each router, or can be a portion, e.g., the low-order bits, of the IP address of the router. Hence, in one embodiment, the hashing operation can be performed on the source ID and destination

ID of the packet being transferred and the IP address of the router. The hash operation produces a result that is unique to each router, even though the operation is performed for the same packet. As a result, the problems described above with conventional hashing procedures are eliminated.

5 In one embodiment, the particular hash operation performed at each node is a cyclic redundancy check (CRC) which includes the additional unique node-identifying input value. The CRC function behaves differently when the node-identifying unique value is used in accordance with the invention. To illustrate, it is assumed that a node is applying a CRC to a set of packets such that there are n possible outcomes. It is further assumed that there is a set of packets such that a CRC applied to those packets will result in only m unique values out of a space of n possible values, where $m < n$. If the CRC is instead applied to the same set of packets using another unique value, i.e., the node-unique value, then the CRC will now result in n unique values in the space of n possible values. This allows an effectively independent hash function to be applied at each node in a network.

15 Specifically, the CRC operation in accordance with the invention has two operands, namely, the input data for which the CRC is to be calculated and a fixed operand, often referred to as the generator polynomial. As known in the art, the generator polynomial is selected based on various characteristics of the CRC calculation, including the size of the CRC result. The CRC calculation is considered to be a division of polynomials with binary coefficients.

20 The CRC calculation involves taking the input data represented by such a polynomial and dividing by the generator polynomial. The remainder of the division is the CRC result. In practice, since a digital computer stores data as binary digits, polynomial arithmetic is implemented as binary arithmetic without carries between bit positions.

5 In one embodiment, the input data to the CRC calculation includes the information used to identify the packet of data being transferred. For example, the input data can include the source address and destination address from the packet header. In addition, the input data to the CRC also includes the node-identifying (node-specific) value. The two addresses as well as the node-specific value can be combined into a single long multiple-bit word. For example, in this illustrative embodiment, the source address and destination address are each 32 bits long, and the node specific value can be 16 bits long. Hence, the input data word to the CRC calculation is a single 80-bit word. This input data word is used to perform the CRC calculation with the preselected generator polynomial to produce the hash result. Because, the node-specific value is used as part of the CRC, the hash value is unique to the particular node router.

10 0964037B-070600
15 In another embodiment, the node-specific value is not part of the input data word to the CRC. Instead, it is used as an initial value for the CRC computation. In the CRC computation, a temporary register is used to store the updated value of the CRC result as the computation is performed. At each step of the calculation, this register is updated with the latest value of the CRC result. In accordance with this embodiment of the invention, the node-specific value is loaded into the temporary register before the calculation begins. The calculation is then performed on the generator polynomial and the input data, which in this case includes only the information from the packet being transferred, for example, the source address and destination address for the packet. Hence, the node-specific value is used in this case as an initial value for the CRC result which is updated as the calculation proceeds. Once again, because the node-specific value is used in the CRC computation, the hash produces a different result at each node for the same data packet.

20
25 While this invention has been particularly shown and described with references to preferred embodiments thereof, it will be understood by those skilled in the art that

various changes in form and details may be made therein without departing from the spirit and scope of the invention as defined by the appended claims.

What is claimed is:

09510278.070600

CLAIMS

1. A method of transferring data on a network, said network including a plurality of nodes connected by a plurality of links, at least one node being connected to more than one link over which the data can be transferred out of the at least one node, the method comprising:
 - providing the data with an identifying portion which identifies a source node and a destination node for the data;
 - associating the at least one node with a node identifying value to distinguish the at least one node from other nodes on the network; and
 - using the identifying portion of the data and the node identifying value, generating a link selection value which identifies one of the more than one links connected to the at least one node to transfer the data out of the at least one node.
2. The method of claim 1 wherein the data is a packet of data.
3. The method of claim 2 wherein the packet is an Internet Protocol (IP) packet.
4. The method of claim 2 wherein the identifying portion of the data is at least a portion of a header of the packet.
5. The method of claim 4 wherein the identifying portion of the data includes a source node ID portion of the header.
6. The method of claim 4 wherein the identifying portion of the data includes a destination node ID portion of the header.

- 1 7. The method of claim 4 wherein the identifying portion of the data includes a
2 protocol portion of the header.
- 1 8. The method of claim 1 wherein the node identifying value is an Internet Protocol
2 (IP) address of the at least one node.
- 1 9. The method of claim 1 wherein generating the link selection value includes
2 performing a logical operation on the identifying portion of the data and the node
3 identifying value.
- 1 10. The method of claim 9 wherein the logical operation comprises a hash operation.
- 1 11. The method of claim 10 wherein the identifying portion of the data includes a
2 source node ID portion of the header.
- 1 12. The method of claim 10 wherein the identifying portion of the data includes a
2 destination node ID portion of the header.
- 1 13. The method of claim 10 wherein the identifying portion of the data includes a
2 protocol portion of the header.
- 1 14. The method of claim 10 wherein the hash operation comprises performing a
2 cyclic redundancy check (CRC) on the identifying portion of the data and the
3 node identifying value.

1 15. An apparatus for transferring data on a network, said network including a
2 plurality of nodes connected by a plurality of links, at least one node being
3 connected to a plurality of links over which the data can be transferred out of the
4 at least one node, the data being provided with an identifying portion which
5 identifies a source node and a destination node for the data, and the at least one
6 node being associated with a node identifying value to distinguish the at least one
7 node from other nodes on the network, the apparatus comprising:

8 an output interface over which the data can be forwarded to one of the
9 plurality of links; and

10 a processing device for generating a link selection value using the
11 identifying portion of the data and the node identifying value, said node selection
12 value identifying one of the plurality of links connected to the at least one node
13 to transfer the data out of the at least one node.

1 16. The apparatus of claim 15 wherein the data transferred by the apparatus is a
2 packet of data.

1 17. The apparatus of claim 16 wherein the packet of data transferred by the
2 apparatus is an Internet Protocol (IP) packet.

1 18. The apparatus of claim 16 wherein the identifying portion of the data used by the
2 processing device to generate the link selection value is at least a portion of a
3 header of the packet.

1 19. The apparatus of claim 18 wherein the identifying portion of the data used by the
2 processing device to generate the link selection value includes a source node ID
3 portion of the header.

- 09640378-070600
- 1 20. The apparatus of claim 18 wherein the identifying portion of the data used by the
2 processing device to generate the link selection value includes a destination node
3 ID portion of the header.
- 1 21. The apparatus of claim 18 wherein the identifying portion of the data used by the
2 processing device to generate the link selection value includes a protocol portion
3 of the header.
- 1 22. The apparatus of claim 15 wherein the node identifying value used by the
2 processing device to generate the link selection value is an Internet Protocol (IP)
3 address of the at least one node.
- 1 23. The apparatus of claim 15 wherein the processing device performs a logical
2 operation on the identifying portion of the data and the node identifying value to
3 generate the link selection value.
- 1 24. The apparatus of claim 23 wherein the logical operation performed by the
2 processing device includes a hash operation.
- 1 25. The apparatus of claim 24 wherein the identifying portion of the data used by the
2 processing device to generate the link selection value includes a source node ID
3 portion of the header.
- 1 26. The apparatus of claim 24 wherein the identifying portion of the data used by the
2 processing device to generate the link selection value includes a destination node
3 ID portion of the header.

1 27. The apparatus of claim 24 wherein the identifying portion of the data used by the
2 processing device to generate the link selection value includes a protocol portion
3 of the header.

1 28. The apparatus of claim 24 wherein the hash operation performed by the
2 processing device comprises a cyclic redundancy check (CRC) on the identifying
3 portion of the data and the node identifying value.

009040" 82201950

APPARATUS AND METHOD FOR EFFICIENT HASHING IN NETWORKS

Abstract of the Disclosure

5 An apparatus and method for efficient hashing uses both an identifying portion of a data packet, e.g., source and destination ID, and an identifying value of the node, e.g., the IP address of the node, to generate a hash result. By inserting a unique value into the hash operation at each node, the invention effectively provides for a different hash implementation at each node. As a result, in situations where multiple paths or multiple links within a path are available to forward packets, traffic can be split over the multiple paths and links. Inefficient utilization of network links found in prior systems which use
10 the same hash operation at each node are eliminated.

K:\IronBridge\5203\patapp2.wpd.doc




Figure 3: A Problem with Conventional Hash Functions



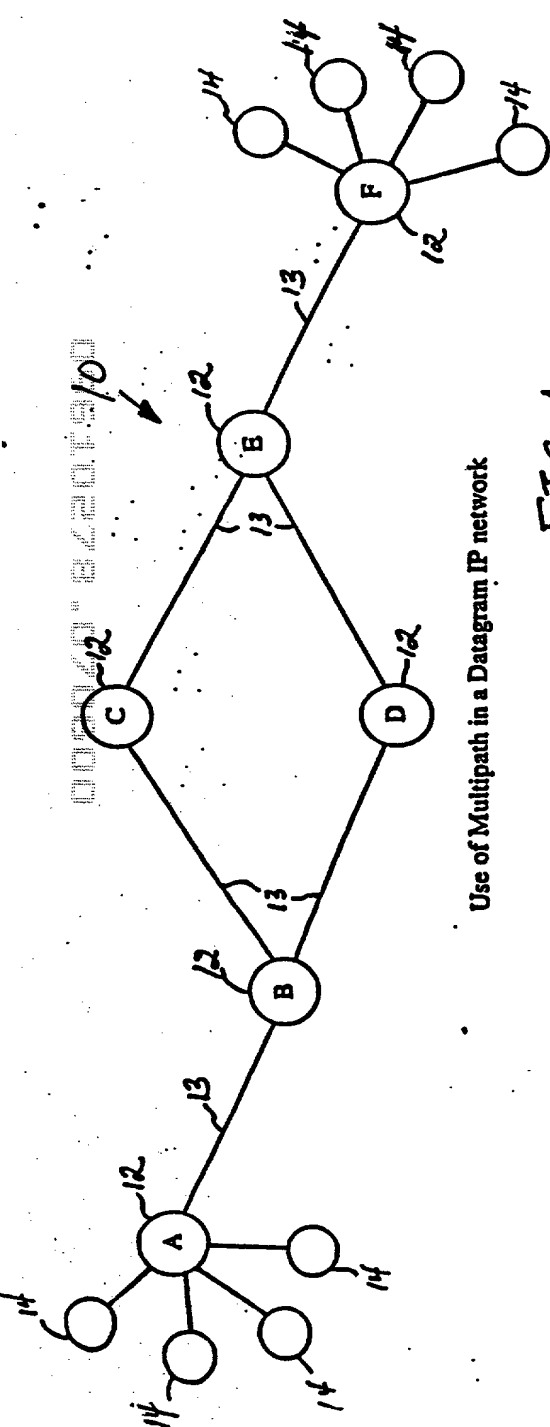


Figure 4: Another Example of The Problem



DECLARATION, PETITION AND POWER OF ATTORNEY FOR PATENT APPLICATION		Attorney Docket No: <div style="text-align: center; font-weight: bold;">IBN.5203</div>	
<p>As a below named inventor, I hereby declare that:</p> <p>My residence, post office address and citizenship are as stated below next to my name,</p> <p>I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:</p> <p>APPARATUS AND METHOD FOR EFFICIENT HASHING IN NETWORKS</p> <p>the specification of which (check only one):</p> <div style="margin-left: 20px;"> <input checked="" type="checkbox"/> is attached hereto. <input type="checkbox"/> was filed as United States Patent Application Serial No. _____ on _____ and was amended on _____ (if applicable) <input type="checkbox"/> was filed as PCT Patent Application Serial No. _____ on _____ and was amended under PCT Article 19 on _____ (if applicable) </div> <p>I hereby state that I have reviewed and understand the contents of the specification, including the claims as amended by any amendment referred to herein.</p> <p>I acknowledge the duty to disclose information which is material to the examination of this application in accordance with Title 37, Code of Federal Regulations, §1.56.</p> <p>I hereby claim foreign priority benefits under Title 35, United States Code, §119 of any foreign application(s) for patent or inventor's certificate or of any PCT international application(s) designating at least one country other than the United States of America listed below and have also identified below any foreign application(s) for patent or inventor's certificate or any PCT international application(s) designating at least one country other than the United States of America filed by me on the same subject matter having a filing date before that of the application(s) of which priority is claimed:</p>			
PRIOR FOREIGN/PCT APPLICATION(S) AND ANY PRIORITY CLAIMS UNDER 35 U.S.C. § 119:			
COUNTRY (if PCT indicate PCT)	APPLICATION NUMBER	DATE OF FILING (day, month, year)	PRIORITY CLAIMED UNDER 35 U.S.C. § 119 (YES/NO)

DECLARATION, PETITION AND POWER OF ATTORNEY FOR PATENT APPLICATION		Attorney Docket No: IBN.5203	
I hereby claim the benefit under Title 35, United States Code, § 119(e) of any United States provisional application(s) listed below.			
PRIOR U.S. APPLICATIONS FOR BENEFIT UNDER 35 U.S.C. § 119(e):			
APPLICATION NUMBER		FILING DATE	
I hereby claim the benefit under Title 35, United States Code, § 120 of any United States application(s) or PCT international application(s) designating the United States of America that is/are listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in that/those prior application(s) in the manner provided by the first paragraph of Title 35, United States Code, § 112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, § 1.56 which occurred between the filing date of the prior applications and the national or PCT international filing date of this application:			
PRIOR U.S. APPLICATIONS OR PCT INTERNATIONAL APPLICATION(S) DESIGNATING THE U.S. FOR BENEFIT UNDER 35 U.S.C. § 120:			
APPLICATION NUMBER (if PCT indicate PCT)		DATE OF FILING (day, month, year)	STATUS: (PATENTED, PENDING OR ABANDONED)
POWER OF ATTORNEY: As a named inventor, I hereby appoint the following attorneys and/or agents to prosecute this application and transact all business in the Patent and Trademark Office connected therewith.			
Maurice E. Gauthier		Reg. No. 20,798	Patrick J. O'Shea
Richard L. Stevens		Reg. No. 24,445	Arlene J. Powers
Matthew E. Connors		Reg. No. 33,298	Steven M. Mills
William E. Hilton		Reg. No. 35,192	Anthony P. Onello, Jr.
			Reg. No. 35,305
			Reg. No. 35,985
			Reg. No. 36,610
			Reg. No. 38,572
Send Correspondence to:		Direct Telephone Calls to:	
Steven M. Mills, Esq. Samuels, Gauthier & Stevens LLP 225 Franklin Street Boston, Massachusetts 02110		Steven M. Mills, Esq. (617) 426-9180 Ext. 149 (617) 426-2275 (facsimile)	
Wherefore I petition that letters patent be granted to me for the invention or discovery described and claimed in the attached specification and claims, and hereby subscribe my name to said specification and claims and to the foregoing declaration, power of attorney, and this petition.			
I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.			
Signature			Date
Full Name of 1st Inventor	Family Name Callon	First Given Name Ross	Second Given Name W.
Residence & Citizenship	City Westford	State or Foreign Country Massachusetts	Country of Citizenship U.S.A.
Post Office Address	Post Office Address 18 Applewood Drive	City Westford	State & Zip Code/Country MA, 01886, U.S.A.

DECLARATION, PETITION AND POWER OF ATTORNEY FOR PATENT APPLICATION			Attorney Docket No: IBN.5203
Signature		Date	
Full Name of 1st Inventor	Family Name Carlson	First Given Name James	Second Given Name D.
Residence & Citizenship	City North Andover	State or Foreign Country Massachusetts	Country of Citizenship U.S.A.
Post Office Address	Post Office Address 25 Essex Street	City North Andover	State & Zip Code/Country MA 01845 U.S.A.
Signature		Date	
Full Name of 1st Inventor	Family Name Ehlert	First Given Name Douglas	Second Given Name C.
Residence & Citizenship	City Dunstable	State or Foreign Country Massachusetts	Country of Citizenship U.S.A.
Post Office Address	Post Office Address 329 Forrest Street	City Dunstable	State & Zip Code/Country MA, 01827, U.S.A.
Signature		Date	
Full Name of 1st Inventor	Family Name Peduel	First Given Name Yuval	Second Given Name
Residence & Citizenship	City Winchester	State or Foreign Country Massachusetts	Country of Citizenship U.S.A.
Post Office Address	Post Office Address 16 Albamont Road	City Winchester	State & Zip Code/Country MA, 01890, U.S.A.

K:\IronBridge\5203\decpoa.wpd